

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Application No. : 10/524,554 Confirmation No. : 9119
First Named Inventor : Richard EVANS
Filed : September 7, 2005
TC/A.U. : 2624
Examiner : Katrina R. Fujita
Docket No. : 038819.55861US
Title : Video Motion Anomaly Detector

DECLARATION OF CLIVE FRENCH UNDER 37 C.F.R. § 1.131

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

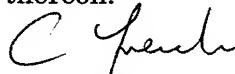
Sir:

I, Clive French, hereby declare that:

1. I am a British patent attorney employed Siemens plc.
2. As part of my duties I prepare patent applications based on materials provided by inventors employed by the Siemens group of companies including Roke Manor Research Limited, the assignee of the present application.
3. Prior to August 11, 2002, Richard John Evans provided me with a copy of his invention disclosure, a copy of which is attached as Exhibit A.
4. Based on Mr. Evans' invention disclosure I prepared and filed a patent application in the British Patent Office. This patent application was accorded British Patent Application No. 0218982.7 and a filing date of August 15, 2002. The present application claims priority from this British patent application.
5. Based on my best recollection, during the period from just prior to August 11, 2002 until August 15, 2002, either I or one of my staff members was diligently preparing the British patent application for filing. The British patent application was mailed to the British Patent Office on August 14, 2002. A copy of the patent application that was mailed to the British Patent Office is attached as Exhibit B.
6. All statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine and/or imprisonment, or both under Section 1001, Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application or any patent issuing thereon.

27 March 2009

Date



Clive French

EXHIBIT A

DRAFT

Video Motion Anomaly Detector (VMAD) – Description

Richard Evans, [REDACTED]

Technical Problem Addressed

The Video Motion Anomaly Detector addresses the problem of automatically detecting events of interest to operators of CCTV systems used in security, transport and other applications, processing CCTV images. The detector may be used in a number of ways, for example to raise an alarm, summoning a human operator to view video data, or to trigger selective recording of video data or to insert an index mark in recordings of video data.

Background to the Problem

Closed circuit television (CCTV) is widely used for security, transport and other purposes. Examples applications include the observation of crime or vandalism in public open spaces or buildings (such as hospitals and school), intrusion into prohibited areas, monitoring the free flow of road traffic, detection of traffic incidents and queues, detection of vehicles travelling the wrong way on one-way roads.

The monitoring of CCTV displays (by human operators) is a very laborious task however and there is considerable risk that events of interest may go unnoticed. This is especially true when operators are required to monitor a number of CCTV camera outputs simultaneously. As a result in many CCTV installations, video data is recorded and only inspected in detail if an event is known to have taken place. Even in these cases, the volume of recorded data may be voluminous and the manual inspection of the data may be laborious. Consequently there is a requirement for automatic devices which process the video images and raise an alarm signal when there is an event of interest. The alarm signal can be used either to draw the event to the immediate attention of an operator, to place an index mark in recorded video or to trigger selective recording of CCTV data.

Some automatic event detectors have been developed for CCTV systems, though few of these are very successful. The most common devices are called video motion detectors (VMDs) or activity detectors, though they are generally based on simple algorithms concerning the detection of changes in the brightness of the video image - not the actual movement of imaged objects. For the purposes of detecting changes in brightness, the video image is generally divided into a grid of typically 16 blocks horizontally and vertically (i.e. 256 blocks in total). There several disadvantages of these algorithms: 1) they are prone to false alarms, for example when there are changes to the overall levels of illumination, 2) they are unable to detect the movement of small objects, because of the block-based processing, 3) they cannot be applied if the scene normally contains movement objects which are not of interest. These disadvantages can be reduced to a limited extent by additional processing logic, but the effectiveness of standard VMDs is inherently limited by the use of change detection as the initial image-processing stage.

There is another type of detection device, which is characterised by the use of complex algorithms involving image segmentation, object recognition and tracking and alarm decision rules. Though these devices can be very effective, they are generally expensive systems designed for use in specific applications and do not

DRAFT

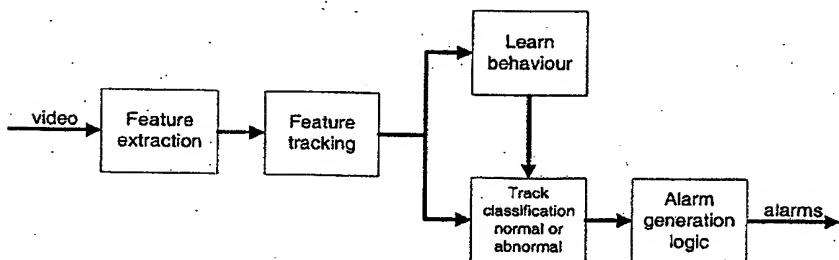
perform well without careful tuning and setting-up, and may not work at all outside of a limited range of applications for which they were originally developed.

As far as is known, the closest thing to the present invention is (in some respects) a device patented by inventors Wade & Jeffrey (Patent No: US6081606, "Apparatus and a method for detecting motion within an image sequence"). It is not known if their invention has been used in any commercial product. Briefly in their invention, motion within the image is calculated by correlating areas of one image with areas of the next image in the video to generate a flow field. The flow field is then analysed and an alarm raised dependent on the observed magnitude and direction of flow. This invention differs significantly from the video motion anomaly detector in that it is not feature based, and alarms are not generated on the basis of abnormal behaviour.

Solution to the problem

The video motion anomaly detector extracts and tracks point-like features in video images and raises an alarm when a feature (or features) is (or are) behaving abnormally compared with the behaviour of features observed over a period of time. By "behaviour" we mean the movement of features in different parts of the video image. For example, rapid movement of features in a particular direction in one part of the field of view may be normal, but it may be abnormal if it occurred in another part of the field of view where the normal behaviour is slow movement. Similarly, rapid movement in the same part of the field of view may be abnormal if the movement is in a different direction.

The following diagram shows the main processing stages in the video motion anomaly detector.



The *feature extraction* stage locates point-like features in each processed image in the video image sequence. A suitable feature has been developed by Harris (Patent No: GB2218507, "Digital Data Processing")

The *feature tracking* stage tracks features so that each point-like feature can be described by its current point and its estimated velocity in the image.

The *learn behaviour* stage accumulates information about the behaviour of features over a period of time. One way of doing this is to accumulate a four-dimensional histogram, the four dimensions of the histogram being x-position, y-position, x-velocity, y-velocity.

DRAFT

The *track classification* stage classifies each track as being normal or abnormal. One way of classifying a track is to compare the frequency of occupancy of the corresponding histogram cell with a threshold. If the frequency of occupancy is below the threshold, the track is classified as abnormal, otherwise it is considered normal.

The *alarm generation* stage generates an alarm signal when abnormal tracks are found to be present, subject to additional processing logic to resolve situations such as intermittent abnormal behaviour or multiple instances of abnormal behaviour associated with one real-world event, and other such situations.

Novel elements of the video motion anomaly detector are

- 1) The use of point feature extraction and tracking in an event detector
- 2) The detection of events by classification of feature behaviour as being abnormal, compared with the behaviour of features observed over time.

Compared with event detection based on normal video motion detection (so called), the video motion anomaly detector has the following advantages.

- 1) It is insensitive to changes in scene illumination levels, which are major source of false alarms in current video motion detectors (because it is based on point feature extraction rather than detecting changes in image brightness).
- 2) It can detect the movement of small objects and raise an alarm if the movement is unusual (because it is based on point features rather than block processing).
- 3) It can detect movements of interest, even in the presence of other objects moving normally (because it accumulates information about feature behaviour).
- 4) It can be applied to a very wide range of different applications with little special setting-up (because it detects abnormal behaviour rather than pre-defined specific behaviour).

Compared with other existing event detection systems based on complex software solutions, the video motion anomaly detector is a simple system suitable for being implemented in inexpensive hardware.

Detailed description of the Invention

... TBC

I would appreciate some feedback on level of detail required etc and other material before drafting this section. – RJE.

DRAFT

Video Motion Anomaly Detector – Abstract

The Video Motion Anomaly Detector addresses the problem of automatically detecting events of interest to operators of CCTV systems used in security, transport and other applications, by processing CCTV images. The detector may be used, for example, to raise an alarm and summon a human operator to view video data, to trigger selective recording of video data or to insert an index mark in recordings of video data. The video motion anomaly detector extracts and tracks point-like features in video images and raises an alarm when a feature (or features) is (or are) behaving abnormally, compared with the behaviour of features observed over a period of time. Compared with existing event detectors called "video motion detectors" (devices which are essentially based on detecting changes in image brightness averaged over image sub-blocks), the video motion anomaly detector has the advantage of being less prone to false alarms caused by changes in scene illumination levels. The video motion anomaly detector can also detect the movement of smaller objects and detect movements of interest in the presence of other moving objects. Further, it can be applied to a very wide range of different applications with little special setting. Compared with other existing event detections systems based on complex software solutions, the video motion anomaly detector can be implemented inexpensive hardware.

Richard Evans



EXHIBIT B

SIEMENS

13 August 2002

JF

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

Our Ref: 2002P10792/GB/R76/CF/GSD

Dear Sirs,

New GB Patent Application
In the name of Roke Manor Research Limited
Entitled "Video Motion Anomaly Detector"

We are filing herewith the following documents:

Enclosures	No. of copies
Patent Form 1/77	1
Patent Form 7/77	2
Fee Sheet	1
Description	2
Drawings	2

If any of the above requires further clarification, please let us know.

Yours faithfully,

C French

Clive French
Intellectual Property Department

enc

Siemens Shared Services Limited

Postal Address
Siemens AG
P.O. Box 22 16 34
80506 Munich
Germany

Office Address
Intellectual Property Department
The Lodge, Roke Manor
Romsey, Hampshire SO51 0ZN

Tel: +44 (1794) 833573
Fax: +44 (1794) 833670

Registered Office: Siemens House, Oldbury, Bracknell, Berkshire, RG12 8FZ Registered No: 2452851

The
Patent
Office

Fee Sheet (FS1)
(Deposit Account Holders)

Account
No. D00393

Payment Method



a) Charge to Deposit Account

b) Cheque enclosed for £ _____

OR

c) Bank transfer of £ _____

Contact Clive French
Tel No. 01794 833573 Date 13 August 2002
Fax No. 01794 833670

Fee Sheet Reference (see Note 3)

2002P10792GB

	Form Type (P/TM/D PCT/EP)	Form Number	Name, Number or Other Identifier	Amount of Fee £
1	P	1/77	2002P10792/GB/R76/CF/GSD	0.00
2	P	7/77	2002P10792/GB/R76/CF/GSD	0.00
3				
4				
5				
6				
7				
8				
9				
10				
11				
12				
13				
14				
15				
16				
17				
18				
19				
20				
21				
22				
23				
24				
25				

Page 1 of 1 pages Total 0.00

Notes for guidance on the completion of this form are printed on the reverse

DDU63/03/99

1/77

Request for grant of a patent*(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)***The Patent Office**Cardiff Road
Newport
Gwent NP10 8QQ

1. Your reference

2002P10792 GB / R76 / CF / GD

2. Patent application number
*(The Patent Office will fill in this part)*3. Full name, address and postcode of the or of
each applicant *(underline all surnames)*ROKE MANOR RESEARCH LIMITED
Old Salisbury Lane, Romsey
Hampshire SO51 0ZNPatents ADP number *(if you know it)*If the applicant is a corporate body, give the
country/state of its incorporation

UNITED KINGDOM

4. Title of the invention

Video Motion Anomaly Detector

5. Name of your agent *(if you have one)**"Address for service" in the United Kingdom
to which all correspondence should be sent
(including the postcode)*Siemens Shared Services Limited
Intellectual Property Department
The Lodge, Roke Manor
Romsey
Hampshire SO51 0ZN
United KingdomPatents ADP number *(if you know it)*

08439358001

6. If you are declaring priority from one or more
earlier patent applications, give the country
and the date of filing of the or of each these
earlier applications and *(if you know it)* the or
each application numberCountry Priority application number
(if you know it) Date of filing
*(day / month / year)*7. If this application is divided or otherwise
derived from an earlier UK application,
give the number and the filling date of
the earlier application8. Is a statement of inventorship and of right
to grant of a patent required in support of
this request? (Answer 'Yes' if:
a) *any application named in part 3 is not an inventor, or*

Yes

b) there is an inventor who is not named as an applicant, or
 c) any named applicant is a corporate body.
 See note (d))

9. Enter the number of sheets for any of the following items you are filling with this form.
 Do not count copies of the same document

Continuation sheets of this form

Description	12
Claim(s)	0
Abstract	0
Drawing(s)	1

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translation of priority documents

Statement of inventorship and right to grant a patent (Patents Form 7/77) 1+1

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11.

I/We request the grant of a patent on the basis of this application

Signature

Date

 Clive French
 Intellectual Property Department 13.08.2002

12. Name and daytime telephone number of Person to contact in the United Kingdom

Clive French

01794 833573

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- a) If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505
- b) Write your answers in capital letters using black ink or you may type them
- c) If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). any continuation sheet should be attached on this form.
- d) If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- e) Once you have filled in the form you must remember to sign and date it.
- f) For details of the fee and ways to pay please contact the Patent Office.

7/77

**Statement of inventorship and of
right to grant of a patent***(See the notes on the back of this form)***The Patent Office**
Cardiff Road
Newport
Gwent NP10 8QQ

1. Your reference

2002P10792 GB / R76 / CF / GD

2. Patent application number
(If you know it)

3. Full name of the or of each applicant

ROKE MANOR RESEARCH LIMITED
Old Salisbury Lane, Romsey
Hampshire SO51 0ZN

4. Title of the invention

Video Motion Anomaly Detector

5. State how the applicant(s) derived the right
from the inventor(s) to be granted a patentBy provisions of contract of service and Section 39(1)
(a) and (b) of the Patent Act 1977.6. How many, if any, additional Patents Forms
7/77 are attached to this form?
(see note (c))

None

7.

I/We believe that the person(s) named over the page
(and on any extra copies of this form) is/are the
inventor(s) of the invention which the above patent
application relates to.

Signature

*C. French*Clive French
Intellectual Property Department

Date

13.08.2002

8. Name and daytime telephone number of
person to contact in the United Kingdom

Clive French

01794 833573

7/77

Enter the full names, addresses and postcodes of the inventors in the boxes and underline the surnames

	1) <u>EVANS</u> Richard John 33 Fordington Ave. Winchester Hampshire SO22 5AN GB
	Patents ADP number (<i>if you know it</i>):
	2)
	Patents ADP number (<i>if you know it</i>):
	3)
	Patents ADP number (<i>if you know it</i>):
4)	
Patents ADP number (<i>if you know it</i>):	
5)	
Patents ADP number (<i>if you know it</i>):	
6)	
Patents ADP number (<i>if you know it</i>):	

7/77

Notes

- a) *If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.*
- b) *Write your answers in capital letters using black ink or you may type them.*
- c) *If there are more than three inventors, please give the names and addresses of the other inventors on the back of another Patents Form 7/77 and attach it to this form.*
- d) *When an application does not declare any priority, or declares priority from an earlier UK application, you must provide enough copies of this form so that the Patent Office can send one to each inventor who is not an applicant.*
- e) *Once you have filled in the form you must remember to sign and date it.*

not the actual movement of imaged objects. For the purposes of detecting changes in brightness, the video image is generally divided into a grid of typically 16 blocks horizontally and vertically (i.e. 256 blocks in total). There are several disadvantages of these algorithms: 1) they are prone to false alarms, for example when there are changes 5 to the overall levels of illumination, 2) they are unable to detect the movement of small objects, because of the block-based processing, 3) they cannot be applied if the scene normally contains movement objects which are not of interest. These disadvantages can be reduced to a limited extent by additional processing logic, but the effectiveness of standard VMDs is inherently limited by the use of change detection as the initial 10 image-processing stage.

There is another type of detection device, which is characterised by the use of complex algorithms involving image segmentation, object recognition and tracking and alarm decision rules. Though these devices can be very effective, they are generally 15 expensive systems designed for use in specific applications and do not perform well without careful tuning and setting-up, and may not work at all outside of a limited range of applications for which they were originally developed.

As far as is known, the closest prior art is represented by a device patented by inventors 20 Wade & Jeffrey (Patent No: US6081606, "Apparatus and a method for detecting motion within an image sequence"). Briefly, motion within the image is calculated by correlating areas of one image with areas of the next image in the video to generate a flow field. The flow field is then analysed and an alarm raised dependent on the observed magnitude and direction of flow. The Wade & Jeffrey invention differs 25 significantly from the video motion anomaly detector of the present invention in that it is not feature based, and alarms are not generated on the basis of abnormal behaviour.

Description of the present invention

30 The video motion anomaly detector extracts and tracks point-like features in video images and raises an alarm when a feature (or features) is (or are) behaving abnormally

compared with the behaviour of features observed over a period of time. By "behaviour" we mean the movement of features in different parts of the video image. For example, rapid movement of features in a particular direction in one part of the field of view may be normal, but it may be abnormal if it occurred in another part of the field of view where the normal behaviour is slow movement. Similarly, rapid movement in the same part of the field of view may be abnormal if the movement is in a different direction.

Fig. 1 shows the main processing stages in the video motion anomaly detector.

10

The feature extraction stage locates point-like features in each processed image in the video image sequence. A suitable feature has been developed by Harris (Patent No: GB2218507, "Digital Data Processing")

15 The feature tracking stage tracks features so that each point-like feature can be described by its current point and its estimated velocity in the image.

The learn behaviour stage accumulates information about the behaviour of features over a period of time. One way of doing this is to accumulate a four-dimensional 20 histogram, the four dimensions of the histogram being x-position, y-position, x-velocity, y-velocity.

25 The track classification stage classifies each track as being normal or abnormal. One way of classifying a track is to compare the frequency of occupancy of the corresponding histogram cell with a threshold. If the frequency of occupancy is below the threshold, the track is classified as abnormal, otherwise it is considered normal.

30 The alarm generation stage generates an alarm signal when abnormal tracks are found to be present, subject to additional processing logic to resolve situations such as intermittent abnormal behaviour or multiple instances of abnormal behaviour associated with one real-world event, and other such situations.

Novel elements of the video motion anomaly detector are

- 1) The use of point feature extraction and tracking in an event detector
- 2) The detection of events by classification of feature behaviour as being abnormal,
5 compared with the behaviour of features observed over time.

Compared with event detection based on normal video motion detection (so called), the video motion anomaly detector has the following advantages.

- 1) It is insensitive to changes in scene illumination levels, which are major source of
10 false alarms in current video motion detectors (because it is based on point feature extraction rather detecting changes in image brightness).
- 2) It can detect the movement of small objects and raise an alarm if the movement is unusual (because it is based on point features rather than block processing).
- 3) It can detect movements of interest, even in the presence of other objects moving
15 normally (because it accumulates information about feature behaviour).
- 4) It can be applied to a very wide range of different applications with little special setting-up (because it detects abnormal behaviour rather than pre-defined specific behaviour).

20 Compared with other existing event detection systems based on complex software solutions, the video motion anomaly detector is a simple system suitable for being implemented in inexpensive hardware.

The main processing stages of the video motion anomaly detector are now described in
25 more detail

1 Feature Extraction

The feature extraction stage locates point-like features in each processed image in the video image sequence. A preferred feature extractor has been developed by Harris
30 (Patent No: GB2218507, "Digital Data Processing") and this is fully described by

Harris and Stephens ("A combined corner and edge detector", Proceedings of the 4th Alvey Vision Conference, September 1988, University of Manchester). An important aspect of this feature extractor is that it provides feature attributes, i.e. quantitative descriptions of the extracted features in addition to their position in the image.

5

The feature extractor may employ a fixed feature-strength threshold, in which case the number of extracted features may vary from frame to frame and it may vary from one application to another. Alternatively the threshold may be varied so that a fixed number of features are extracted.

10

Compared with other point-feature extraction algorithms the Harris method is particularly robust. Other point-feature extraction techniques have been developed, for example by Moravec ("Obstacle avoidance and navigation in the real world by a seeing robot rover" Tech Report CMU-RI-TR-3, Carnegie-Mellon University, Robotics Institute, Sep 1980), and other ad hoc schemes can be envisaged. Indeed any algorithm, extracting image features that can be associated with a locality, can be used as a point-like feature extractor. As examples, knot-points (i.e. points of high curvature on edge features) can be assigned a position, and an image region (perhaps an entire vehicle or person, segmented by edge-finding or region growing techniques) can be 20 assigned the position of its centroid.

2 Feature Tracking

The feature tracking stage tracks features between image frames so that each point-like 25 feature can be described by its current position and its estimated velocity in the image.

Tracking algorithms are a relatively well-known and understood. A treatise on the subject is given by Blackman & Popoli ("Design & Analysis of Modern Tracking Systems, Artech House 1999). In the present application a multi-target tracking 30 algorithm is required as a scene may contain a large number of moving objects and each may give rise to a number of extracted features. For example, a car passing

through the field of view may generate a number of extracted features on each video frame, depending on the spatial resolution used, and a traffic scene may contain a number of cars moving in different parts of the field of view.

5 Tracking is an iterative process. At any time, a number of objects may be being tracked; i.e. their current position and velocity are known and as each new frame of video data is presented, the track information needs to be updated. Such a tracking algorithm typically consists of the following stages:

10. Plot (feature)-to-track Association: this is the process of deciding which of the features, extracted from the most recent video frame, is the same object as any particular track. As "plot" rather than "feature" is the normal term used in discussion of tracking algorithms, "plot" will be used in this section of the algorithm description. A standard approach is to only consider association plots that fall within a window or gate, 15 which is centred on the plot's predicted position (see below). Many tracking algorithms then employ simple rules to handle situations where more than one plot falls in the acceptance gate. Example rules include 'associate the plot nearest to the predicted position', or 'associate neither' (see below). In the present application the density of plots is typically high and the possibility of plot-to-track association error is 20 high so the preferred approach makes use of the similarity of plot (feature) attributes as well as plot position in making the plot-to-track association decision. Other schemes for resolving ambiguities are possible, for example probabilistic matching, multiple-hypothesis tracking and deferred decision making. Other refinements may be employed to improve performance, for example cross-correlation of image patches to 25 confirm the similarity of imaged features, variation of the acceptance window size according to the expected accuracy of prediction (thus well established and slow moving tracks may have a smaller acceptance window than fast moving or newly formed tracks), and bi-directional matching schemes.
- 30 Track Maintenance: this is the process of initiating new tracks (because a new object has come into view) and deleting tracks (because the tracked feature is no longer in

view); most tracking algorithms will also have a track confirmation process (i.e. deciding a track is sufficiently well-established as to be the basis for subsequent decision making). In the preferred implementation, new tracks are initiated from plots that cannot be accounted for after plot-to-track association, i.e. association of plots with 5 existing tracks. Because of uncertainties in the performance of the feature extractor, tracks are not immediately deleted in the preferred implementation if they are not associated with any plot; instead a track may "coast" for a number of frames before deletion. Tracks are confirmed once they have been successfully associated with plots on a number of frames.

10

Tracking features which have low feature strength is problematic, because they are not reliably detected on each video frame, and tracking errors become more common when the tracks are closely spaced. One way of reducing this problem is to ignore unmatched plots of a low feature-strength, and only initiate new tracks for unmatched 15 plots of a higher feature-strength.

Track Filtering and Prediction: this is the process of estimating the current plot position and velocity, and predicting the expect plot position on subsequent image frames. This process is required because measurements of feature positions may be imperfect 20 because of pixel quantisation and image noise, and objects may move appreciably between image frames. A number of methods are applicable, for example recursive Kalman or Alpha-Beta filtering, fitting polynomial splines to recent plot data etc. Performance here may be improved by a number of schemes, for example outlier removal and varying the order of a polynomial spline according the length of a track's 25 history.

In general tracking algorithms follow a common pattern though individual implementations may vary in detail according to application-specific factors and other issues.

30

3 Learn Behaviour

The learn behaviour stage accumulates information about the behaviour of features over a period of time. The preferred way of doing this is to accumulate a four-dimensional histogram, the four dimensions of the histogram being x-position, y-position, x-velocity, y-velocity.

The behaviour histogram, being four-dimensional, may require a large amount of data storage and this may be impractical for implementation in processing hardware. The preferred way of overcoming this problem is to partition the histogram into two sections. The first is for stationary tracks, i.e. very slow moving ones. This section is a two-dimensional histogram and so may use a fine cell size without requiring large amounts of storage. The second histogram section is for moving objects of different speeds. Although this is a four-dimensional histogram, a coarser cell size may be used for such objects so the size of this section of the histogram also need not be large. The size of the second partition may also be reduced by using a cell size which varies with speed.

There are a number of alternative ways of constructing the histogram to minimise memory requirements. These include quad tree and other sparse matrix representation techniques.

The histogram describing behaviour may be built up over a fixed period of time and remain unchanged thereafter. This has the disadvantage that slow changes in actual behaviour, drift in camera electronics or sagging of camera mounts etc, may ultimately result in normal behaviour being classified as abnormal. This may be overcome in a number of ways. In the preferred method, the histogram is periodically de-weighted by a factor close to unity, with the result that the system has a fading memory, i.e. its memory is biased towards most recent events. Depending on processor limitations, it may be necessary to implement the fading memory in ways to reduce processor load.

For example, the de-weighting might be applied at a larger interval, or only a part of it

might be de-weighted at shorter intervals. Another way of biasing memory of behaviour to recent events, albeit not a fading memory, is to use two histogram stores, one being built up while the other is being used for track classification.

5 While histogramming is the preferred approach, there are other possible approaches and these are discussed below under the heading "Alternative Learning and Classification Methods".

4 Track Classification

10

The track classification stage classifies each track as being normal or abnormal and the method used depends on how behaviour is being learnt. In the preferred histogram-based method, a track is classified by comparing the frequency of occupancy of the corresponding histogram cell with an occupancy threshold. If the frequency of 15 occupancy is below the threshold, the track is classified as abnormal, otherwise it is considered normal.

If a very low false alarm rate is required, it may be necessary to take additional steps to prevent an adverse system response, particularly if the training time has been limited.

20 As an example of such steps, a track may be classified as normal, even if the occupancy of the corresponding histogram cell is low, if it is adjacent or near an above threshold cell. The distance from the corresponding histogram cell to the nearest above occupancy threshold cell (measured within the histogram by city-block, Cartesian or some other metric which may be occupancy weighted) may be compared with a 25 distance threshold. If the distance is less than the threshold, the track is classified as normal, otherwise the track is considered abnormal. The false alarm rate can be adjusted by varying the occupancy threshold and the distance threshold. As an alternative to the use of a distance threshold, histogram data might be blurred to suppress the classification of tracks as abnormal in cells close to high occupancy cells.

30 Similarly, other morphological operators might be applied.

5 Alarm Generation

The alarm generation stage generates an alarm signal when abnormal tracks are found to be present, subject to additional processing logic to resolve situations such as 5 intermittent abnormal behaviour or multiple instances of abnormal behaviour associated with one real-world event, and other situations such as when spurious track data is generated by a tracking error.

The risk of alarms being generated by spurious data can be reduced by limiting 10 processing to confirmed tracks or tracks whose track history shows a consistent history of associated plots.

To prevent intermittent alarms, an alarm can be raised (subject to other logic) only 15 when a track is classified as abnormal for the first time, or when a filtered version of the classification rises above a threshold.

Multiple alarms might be generated, for example, if a vehicle viewed by the CCTV system takes an abnormal path and generates a number of separate abnormal tracks in the process – the different tracks being generated by different parts of the vehicle. 20 These multiple alarms would be confusing and unwanted in a practical system. These can be suppressed by inhibiting alarms for a period of seconds after a first alarm. Alternatively alarms could be suppressed if the track causing the alarm is within some distance of a track that has previously generated an alarm.

25 A number of different methods of alarm generation logic can be envisaged, with different ad hoc formulations.

6 Alternative Learning and Classification Methods

30 The above sections describe preferred methods though there are alternatives and a selection is described below.

6.1 Fan-in/Fan-out (MLP) Neural Net

In outline, the idea is to use a multi-layer perceptron with 4 input nodes (for track data X, Y, Vx, Vy) and 4 output nodes, but an intermediate layer of 2 or 3 nodes. The network would be trained to reconstruct its own input. Because of the internal constriction, the reconstruction should be better for "normal" tracks and worse for "abnormal" tracks. Thus, the accuracy of reconstruction could be used to assess the normality of the track.

10

6.2 Nearest Neighbour

Here tracked feature data for recent frames is retained to create a track history database. Each new track is tested for normality by searching this database to find any similar 15 previous tracks.

6.3 Pruned Nearest Neighbour

This is a variation of the full nearest neighbour technique. The history database is 20 reduced in size by omitting duplicates or near duplicates of earlier data.

6.4 Kohonen Net

Though a neural net technique, this can be viewed as similar to the pruned nearest 25 neighbour method. The behaviour of tracked objects is described by a set of nodes positioned in the four-dimensional input space. The actual positions of the nodes are determined by an iterative training process. This method is also related to adaptive code-book generation methods used in data compression systems.

30 6.5 Probabilistic Checking

This is a lateral approach to searching history databases for the nearest neighbour-based algorithms. Here, the history database is searched by choosing for comparison entries in a random sequence until a number of matches are found. If the track being checked is very normal, a number of matches will be found very quickly.

5

Accordingly, the present invention provides a Video Motion Anomaly Detector which addresses the problem of automatically detecting events of interest to operators of CCTV systems used in security, transport and other applications, by processing CCTV images. The detector may be used, for example, to raise an alarm and summon a 10 human operator to view video data, to trigger selective recording of video data or to insert an index mark in recordings of video data. The video motion anomaly detector extracts and tracks point-like features in video images and raises an alarm when a feature (or features) is (or are) behaving abnormally, compared with the behaviour of features observed over a period of time. Compared with existing event detectors called 15 "video motion detectors" (devices which are essentially based on detecting changes in image brightness averaged over image sub-blocks), the video motion anomaly detector has the advantage of being less prone to false alarms caused by changes in scene illumination levels. The video motion anomaly detector can also detect the movement of smaller objects and detect movements of interest in the presence of other moving 20 objects. Further, it can be applied to a very wide range of different applications with little special setting. Compared with other existing event detection systems based on complex software solutions, the video motion anomaly detector can be implemented inexpensive hardware.

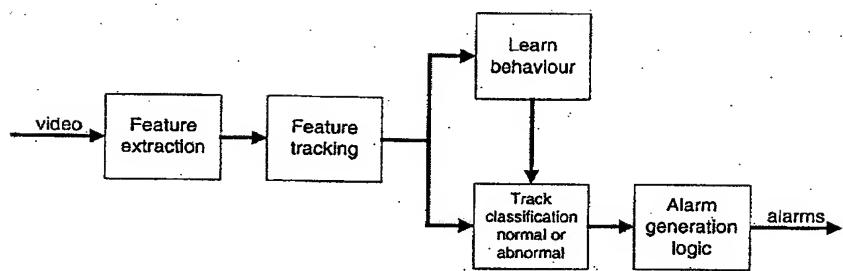


Fig. 1